

What is claimed is:

1. A data protection system comprising:
 - a fileserver having:
 - a filter driver operative to intercept input/output activity initiated by client file requests and to maintain a list of modified and created files since a prior backup;
 - a file system in communication with the filter driver and operative to store client files;
 - a policy cache operative to store a protection policy associated with a set of files;
 - a mirror service in communication with the filter driver and with the policy cache, the mirror service operative to prepare modified and created files in the set of files to be written to a repository as specified in the protection policy associated with the set of files;
 - a fileserver API coupled to the mirror service and operative to communicate with a repository; and
 - a fileserver file transfer module in communication with the file system and operative to transfer files from the file system to at least one repository.
2. The system of claim 1 wherein the mirror service directs new versions of an existing file to the repository to which prior versions of the file were written.

3. The system of claim 1 wherein the system further comprises:

a location cache in communication with the mirror service and operative to indicate which repository should receive an updated version of an existing file; and

a location manager coupled to the location cache and operative to update the location cache when the system writes a new file to a specific repository node.

4. The system of claim 3 wherein the system further comprises

a local repository having:

a local repository node API adapted for communicating with the fileserver API;

a local repository file transfer module in communication with the fileserver file transfer module and adapted for receiving files from the fileserver file transfer module;

a data mover in communication with the local repository API and operative to supervise the replication of files from the fileserver to the local repository; and

a protection policy component in communication with the data mover and operative to determine whether new versions of existing files should be compressed and whether older versions of existing files should be maintained.

5. The system of claim 4 wherein the system further comprises:

a remote repository having:

a remote repository node API adapted for communicating with the local repository API;

a remote repository file transfer module in communication with the local file transfer module and adapted for receiving files from the local file transfer module;

 a data mover in communication with the remote repository API and operative to supervise the replication of files from the local repository to the remote repository; and

 a protection policy component in communication with the data mover and operative to determine whether new versions of existing files should be compressed and whether older versions of existing files should be maintained.

6. The system of claim 1 wherein the protection cache is operative to define which repositories are used, how often data protection occurs, how many replicas are maintained within each repository, and how modifications to share data are maintained.

7. A method for protecting data comprising:

 storing a version of a first file within a set of files on a primary disk storage system;

 examining a protection policy associated with the set of files to determine where and how to protect files associated with the set of files; and

 replicating the version of the first file to repositories specified by the protection policy, the specified repositories including at least one local repository and at least one remote repository.

8. The method of claim 7 wherein the version of the first file is the first version.

9. The method of claim 8 wherein the method further comprises:
 - applying reverse delta compression to successive versions of the first file as new versions are stored in the repositories.
10. The method of claim 9 wherein applying reverse delta compression to successive version of the first file comprises in response to the creation of a second version of the first file:
 - replacing the first version of the first file replicated in the local repository with a reverse delta compressed version representing the difference between the first version and the second version and replicating the second version in the local repository;
 - transmitting a difference file to the remote repository; and
 - in the remote repository, applying the difference file to the previous version of the file to store the second version and a reverse delta compressed version representing the difference between the first version and the second version.
11. The method of claim 7 wherein examining a protection policy associated with the share to determine where and how to protect files associated with the set of files comprises:
 - determining the location of repositories and the number of replicas for each repository.
12. The method of claim 7 wherein examining a protection policy associated with the set of files to determine where and how to protect files associated with the set of files comprises:
 - determining whether to purge a file from repositories after the file has been deleted from a set of files.

13. The method of claim 7 wherein examining a protection policy associated with the set of files to determine where and how to protect files associated with the set of files comprises:

determining whether to keep version histories.

14. The method of claim 7 wherein examining a protection policy associated with the set of files to determine where and how to protect files associated with the set of files comprises:

determining a specified backup frequency.

15. The method of claim 7 wherein examining a protection policy associated with the set of files to determine where and how to protect files associated with the set of files comprises:

determining a specified type of compression.

16. The method of claim 7 wherein examining a protection policy associated with the set of files to determine where and how to protect files associated with the set of files comprises:

determining a specified caching level.

17. A data protection system comprising:

a fileserver having:

filter driver means for intercepting input/output activity initiated by client file requests and for maintaining a list of modified and created files since a prior backup;

file system means in communication with the filter driver, the file system means for storing client files;

policy cache means for storing a protection policy associated with a set of files;

mirror service means in communication with the filter driver means and with the policy cache means, the mirror service means for preparing modified and created files in the set of files to be written to a repository as specified in the protection policy associated with the set of files.

18. The system of claim 17 wherein the system further comprises:

a fileserver API coupled to the mirror service means and operative to communicate with a repository; and

a fileserver file transfer module in communication with the file system means and operative to transfer files from the file system to at least one repository.